



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

am

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/872,891	05/31/2001	Sashikanth Chandrasekaran	256/145	3158
23639	7590	06/16/2005	EXAMINER	
BINGHAM, MCCUTCHEN LLP THREE EMBARCADERO CENTER 18 FLOOR SAN FRANCISCO, CA 94111-4067			GOLD, AVI M	
			ART UNIT	PAPER NUMBER
			2157	

DATE MAILED: 06/16/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

09/872,891

Applicant(s)

CHANDRASEKARAN ET AL.

Examiner

Avi Gold

Art Unit

2157

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 02 May 2005.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-66 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-66 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____.
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____.
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____.

DETAILED ACTION

This action is responsive to the amendment filed on May 2, 2005. Claims 1, 14, 32-36, 40, and 54 were amended. Claims 1-66 are pending.

Response to Amendment

Claim Rejections - 35 USC § 102

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

2. Claims 1, 14, 32-36, 40, and 54 are rejected under 35 U.S.C. 102(b) as being anticipated by Rao et al., U.S. Patent No. 5,657,450.

Rao teaches the invention as claimed including accessing information from distal sources on a network (see abstract).

Regarding claims 1, 36, and 40, Rao teaches a method, computer program product, and system for predicting the behavior of a workload across a plurality of nodes, comprising:

a) receiving a workload to be executed (col. 2, lines 36-40, Rao discloses an access operation received on an intermediary server);

b) executing the workload on a single node before the workload is sent to a plurality of nodes for execution (col. 2, lines 40-45, Rao discloses analyzation of the operation before it is sent anywhere);

c) tracing the execution of the workload to identify a potential data conflict (col. 2, lines 4-13, 28-35, and 45-50, Rao discloses the synthesizing of time and progress estimates based on possible paths for the operation to take, along the way taking various factors into account);

d) based on a result of the tracing, predicting the behavior of the workload across the plurality of nodes (col. 2, lines 4-13, 28-35, and 45-53); and

e) outputting the prediction (col. 2, lines 4-13, 28-35, and 45-53, Rao discloses providing the time estimate to the user).

Regarding claims 14, 33, 34, and 54, Rao teaches a method, computer program product, and system for distributing a workload across a plurality of nodes, the method comprising:

a) receiving a workload to be executed;

b) executing the workload on a single node before the workload is sent to a plurality of nodes for execution;

c) tracing the execution of the workload to identify a potential data conflict;

d) based on a result of the tracing, forming a workload distribution scheme that distributes the workload across the plurality of nodes; and

e) outputting the workload distribution scheme (col. 2, lines 1-53).

Regarding claims 32 and 35, Rao teaches a computer program product that includes a medium usable by a processor, the medium comprising a sequence of instructions which, when executed by said processor, causes said processor to execute a process for optimizing the distribution of a workload across a plurality of nodes, the process and system comprising:

- a) receiving a workload to be executed;
- b) executing the workload on a single node before the workload is sent to a plurality of nodes for execution;
- c) tracing the execution of the workload to identify a potential data conflict;
- d) based on a result of the tracing, optimizing the distribution of the workload across the plurality of nodes; and
- e) outputting the optimized distribution scheme (col. 2, lines 1-53).

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

4. Claims 1-9, 11, 14-58, 60-63, 65, and 66 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gersht et al., U.S. Patent No. 6,405,257, further in view of Rao et al., U.S. Patent No. 5,657,450.

Gersht teaches the invention substantially as claimed including methods and systems for congestion avoidance in IP backbone networks (see abstract).

Regarding claims 1, 36, and 40, Gersht teaches a method, computer program product, and system for predicting the behavior of a workload across a plurality of nodes (col. 2, lines 10-28, 34-41), comprising:

a) receiving a workload to be executed (col. 2, lines 10-13, Gersht discloses traffic directed to a source node);

b) executing the workload on a single node before the workload is sent to a plurality of nodes for execution (col. 2, lines 10-16);

c) tracing the execution of the workload to identify a potential data conflict (col. 2, lines 15-20, Gersht discloses a predetermined sets of routes for traffic);

d) based on a result of the tracing, predicting the behavior of the workload across the plurality of nodes (col. 2, lines 15-20, Gersht discloses a predetermined sets of routes for traffic based on different factors); and

e) outputting the prediction (col. 2, lines 34-41, Gersht discloses network computers maximum permitted rate for each predetermined route in the network).

Gersht fails to teach the limitation further including executing the workload on a single node before it is sent to a plurality of nodes for execution.

However, Rao teaches accessing information from distal sources on a network (see abstract). Rao teaches the use of analyzation of an operation on a node before it is sent anywhere (col. 2, lines 40-45).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Gersht in view of Rao to execute the workload on a single node before it is sent to a plurality of nodes for execution. One would be motivated to do so because it is beneficial for a user to know how much time is required for performing an operation on a distal source (Rao, col. 1, lines 19-22).

Regarding claims 2, 37, 41, 55, 60, and 65, Gersht teaches the method, computer program product, and system of claims 1, 36, 40, 54, 32, and 35 wherein the action of identifying potential data conflicts comprises predicting how many data conflicts will occur (col. 2, lines 15-20, Gersht discloses checks on peak packet rate, service class, and quality of service requirements).

Regarding claims 3, 38, 42, 56, 61, and 66, Gersht teaches the method, computer program product, and system of claims 1, 36, 40, 54, 32, and 35 wherein the action of identifying potential data conflicts comprise predicting types of data conflicts (col. 2, lines 15-20).

Regarding claim 4, Gersht teaches the method of claim 3 in which the types of data conflicts comprises a read-write conflict (col. 3, lines 21-25, Gersht discloses adjustment of routes based on changes to traffic demand).

Regarding claim 5, Gersht teaches the method of claim 3 in which the types of data conflicts are based upon types of operations needed to resolve the data conflicts (col. 2, lines 10-14, Gersht discloses burst access thresholds).

Regarding claim 6, Gersht teaches the method of claim 3 in which the different types of data conflicts have differing levels of expense associated with operations needed for data conflict resolution (col. 2, lines 15-20, Gersht discloses QoS requirements).

Regarding claims 7, 57, and 62, Gersht teaches the method and computer program product of claims 1, 54, and 32 in which the potential data conflicts are at the granularity of a data block (col. 2, lines 15-20, Gersht discloses peak packet rate).

Regarding claims 8, 39, 43, 58, and 63, Gersht teaches the method, system, and computer program product of claims 1, 36, 40, 54, and 32 in which the potential data conflicts are identified based upon workload division between sessions (col. 3, lines 21-25, Gersht discloses a change of traffic demands in intervals).

Regarding claim 9, Gersht teaches the method of claim 1 further comprising:

f) selecting a number of nodes (col. 2, lines 15-20, Gersht discloses predetermined routes along nodes);

g) dividing the traced execution of the workload across the number of nodes (col. 2, lines 10-14, Gersht discloses a preallocated set of maximum permitted rates and preassigned set of burst access thresholds).

Regarding claim 11, Gersht teaches the method of claim 9 in which the number of nodes corresponds to an anticipated number of nodes for a distributed computing system (col. 4, lines 14-35, Gersht discloses nodes with a certain function).

Regarding claims 14, 33, 34, and 54, Gersht teaches a method, computer program product, and system for distributing a workload across a plurality of nodes, the method comprising:

- a) receiving a workload to be executed;
- b) executing the workload on a single node before the workload is sent to a plurality of nodes for execution;
- c) tracing the execution of the workload to identify a potential data conflict;
- d) based on a result of the tracing, forming a workload distribution scheme that distributes the workload across the plurality of nodes; and
- e) outputting the workload distribution scheme (col. 2, lines 10-28).

Gersht fails to teach the limitation further including executing the workload on a single node before it is sent to a plurality of nodes for execution.

However, Rao teaches the use of analyzation of an operation on a node before it is sent anywhere (col. 2, lines 40-45).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Gersht in view of Rao to execute the workload on a single node before it is sent to a plurality of nodes for execution. One would be motivated to do so because it is beneficial for a user to know how much time is required for performing an operation on a distal source (Rao, col. 1, lines 19-22).

Regarding claims 15, 44, and 49, Gersht teaches the method, computer program product, and system of claims 14, 33, and 34, wherein the forming the workload distribution scheme comprises determining workload distribution in manner which reduces the potential data conflicts (col. 2, lines 15-20).

Regarding claims 16, 45, and 50, Gersht teaches the method, computer program product, and system of claims 14, 33, and 34, wherein the workload distribution scheme is based upon data accesses (col. 2, lines 20-28, Gersht discloses burst access).

Regarding claim 17, Gersht teaches the method of claim 16 in which the workload is grouped in the workload distribution scheme to maximize intersection of

data access on a same group of nodes (col. 2, lines 10-14, Gersht discloses maximum rates).

Regarding claim 18, Gersht teaches the method of claim 16 in which the workload is grouped in the workload distribution scheme to minimize intersection of data access across different groups of nodes (col. 3, lines 15-20, Gersht discloses routes between different source-destination node pairs).

Regarding claims 19, 46, and 51, Gersht teaches the method, computer program product, and system of claims 14, 33, and 34, wherein the workload distribution scheme is based upon access frequencies (col. 2, lines 15-20, Gersht discloses a peak packet rate).

Regarding claim 20, Gersht teaches the method of claim 19 in which data objects accessed by the workload are associated with weighting factors (col. 2, lines 15-20).

Regarding claim 21, Gersht teaches the method of claim 20 in which not all the data objects are associated with same weighting factors (col. 2, lines 15-20, Gersht discloses a packet rate and a QoS requirement).

Regarding claim 22, Gersht teaches the method of claim 20 in which a weighted correlation is performed between the data objects and entities that access the data objects (col. 2, lines 15-20).

Regarding claim 23, Gersht teaches the method of claim 22 in which the entities that access the data objects comprises sessions (col. 2, lines 15-20).

Regarding claim 24, Gersht teaches the method of claim 22 in which subsets of the entities that access the data objects are grouped together (col. 2, lines 15-28).

Regarding claim 25, Gersht teaches the method of claim 24 in which a data structure is employed to represent an affinity between one of the entities that access the data objects and another of the entities (col. 2, lines 10-28, Gersht discloses a path based on all nodes).

Regarding claims 26, 47, and 52, Gersht teaches the method, computer program product, and system of claims 14, 33, and 34 in which the workload comprises data access upon one or more hierarchical objects (col. 2, lines 10-28, Gersht discloses nodes in a route).

Regarding claim 27, Gersht teaches the method of claim 26 in which tracing the execution of the workload comprises tracing identifiers for the one or more hierarchical objects (col. 2, lines 20-28, Gersht discloses identifying a route for the burst).

Regarding claims 28, 48, and 53, Gersht teaches the method, computer program product, and system of claims 14, 33, and 34 in which tracing the execution of the workload comprises tracing identifiers associated with entities that access data (col. 3, lines 29-39, Gersht discloses a burst controller).

Regarding claim 29, Gersht teaches the method of claim 28 in which the entities comprise sessions (col. 3, lines 29-39, Gersht discloses bursts).

Regarding claim 30, Gersht teaches the method of claim 28 in which the workload distribution scheme distributes the workload based upon partitioning of the entities that access data (col. 3, lines 29-39, Gersht discloses predetermined route for the detected burst based on certain factors).

Regarding claim 31, Gersht teaches the method of claim 30 in which an association is formed between partitioning of the entities that access data and partitioning of one or more applications within the workload (col. 3, lines 29-39, Gersht discloses a predetermined route, peak packet rate, and burst access threshold all being factors).

Regarding claims 32 and 35, Gersht teaches a computer program product that includes a medium usable by a processor, the medium comprising a sequence of instructions which, when executed by said processor, causes said processor to execute a process for optimizing the distribution of a workload across a plurality of nodes, the process and system comprising:

- a) receiving a workload to be executed;
- b) executing the workload on a single node before the workload is sent to a plurality of nodes for execution;
- c) tracing the execution of the workload to identify a potential data conflict;
- d) based on a result of the tracing, optimizing the distribution of the workload across the plurality of nodes; and
- e) outputting the optimized distribution scheme (col. 2, lines 10-28).

Gersht fails to teach the limitation further including executing the workload on a single node before it is sent to a plurality of nodes for execution.

However, Rao teaches the use of analyzation of an operation on a node before it is sent anywhere (col. 2, lines 40-45).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Gersht in view of Rao to execute the workload on a single node before it is sent to a plurality of nodes for execution. One would be motivated to do so

because it is beneficial for a user to know how much time is required for performing an operation on a distal source (Rao, col. 1, lines 19-22).

5. Claims 10 and 12 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gersht further in view of Martin et al., U.S. Patent No. 6,154,813.

Gersht teaches the invention substantially as claimed including methods and systems for congestion avoidance in IP backbone networks (see abstract).

As to claims 10 and 12, Gersht teaches the method of claim 9.

Gersht fails to teach the limitation further including the use of modulo division to divide the traced execution of the workload across the number of nodes and the use of a modulo class to represent a node in the number of nodes.

However, Martin teaches a cache management scheme for continuous media data, such as audio or video (see abstract). Martin teaches the use of modulo division (col. 4, lines 1-15, col. 5, lines 46-63).

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Gersht in view of Martin to use modulo division and a modulo class in association with nodes. One would be motivated to do so because they are efficient ways of organizing nodes.

6. Claims 13, 59, and 64 are rejected under 35 U.S.C. 103(a) as being unpatentable over Gersht further in view of Auvenshine, U.S. Patent No. 6,542,930.

Gersht teaches the invention substantially as claimed including methods and systems for congestion avoidance in IP backbone networks (see abstract).

As to claims 13, 59, and 64, Gersht teaches the method and computer program product of claims 1, 54, and 32.

Gersht fails to teach the limitation further including the potential data conflicts being used to compute costs of migrating the workload to a distributed system.

However, Auvenshine teaches a distributed file system with automated file management achieved by decoupling data analysis and movement operations (see abstract). Auvenshine teaches the use of a distributed system.

It would have been obvious to one of ordinary skill in the art at the time of the invention to modify Gersht in view of Auvenshine to migrate the workload to a distributed system. One would be motivated to do so because it would still seem as if the system is one local machine.

Response to Arguments

7. Applicant's arguments with respect to claims 1-66 have been considered but are moot in view of the new ground(s) of rejection.

Conclusion

8. The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

U.S. Pat. No. 6,427,166 to Hurst et al.

U.S. Pat. No. 5,928,344 to Stierli.

U.S. Pat. No. 6,681,251 to Leymann et al.

U.S. Pat. No. 6,442,564 to Frey et al.

U.S. Pat. No. 5,819,033 to Caccavale.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Avi Gold whose telephone number is 571-272-4002. The examiner can normally be reached on M-F 8:00-5:30 (1st Friday Off).

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Ario Etienne can be reached on 571-272-4001. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Avi Gold

Patent Examiner

Art Unit 2157

AMG


ARIO ETIENNE
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100